# Artificial Intelligence Explained

Author, Francesco Corea

A Data Science Foundation White Paper

January 2018

---

www.datascience.foundation

# Artificial Intelligence Explained

## Basic definitions and categorization

### *I. Overview*

Artificial Intelligence (AI) represents nowadays a paradigm shift that is driving at the same time the scientific progress as well as the industry evolution. Given the intense level of domain knowledge required to really appreciate the technicalities of the artificial engines, what AI is and can do is often misunderstood: the general audience is fascinated by its development and frightened by terminator-like scenarios; investors are mobilizing huge amounts of capital but they have not a clear picture of the competitive drivers that characterize companies and products; and managers are rushing to get their hands on the last software that may improve their productivities and revenues, and eventually their bonuses.

Even though the general optimism around creating advancements in artificial intelligence is evident (Muller and Bostrom, 2016), in order to foster the pace of growth facilitated by AI I believe it would be necessary to clarify some concepts.

### *II. Basic Definitions Categorization*

First, let's describe what artificial intelligence means. According to Bostrom (2014), AI today is perceived in three different ways: it is something thatmight answer all your questions, with an increasing degree of accuracy ("the Oracle"); it could do anything it is commanded to do ("the Genie"), or it might act autonomously to pursue a certain long-term goal ("the Sovereign"). However, AI should not be defined by what it can do or not, and thus a broader definition is appropriate.

*An artificial intelligence is a system that can learn how to learn, or in other words a series of instructions (an algorithm) that allows computers to write their own algorithms without being explicitly programmed for.*

Although we usually think about intelligence as the computational part of our ability to achieve certain goals, it is rather the capacity to learn and solve new problems in a changing environment. In a primordial world then, it is simply the attitude to foster survival and reproduction (Lo, 2012; 2013; Brennan and Lo, 2011; 2012). A living being is then defined as intelligent if she is driving the world into states she is optimizing for.

No matter how accurately we defined this concept, we can intuitively understand that the level of intelligence machines are provided with today is years far from the average level of any human being. While human being actions proceed from observing the physical world and deriving underlying relationships that link cause and effect in natural phenomena, an artificial intelligence is moved entirely by data and has no prior knowledge of the nature of the relationship among those data. It is then "artificial" in this sense because it does not stem from the physical law but rather from pure data.

We then have just defined what artificial intelligence is and what mean to us. In addition to that, though, there are two other concepts that should be treated as part of this introduction to AI: first of all, how AI is different and/or related to other buzzwords (big data, machine learning, etc.); second, what features a system has to own to be defined intelligent.

We think of AI as an interdisciplinary field, which covers (and requires) the study of manifold sub-disciplines, such as natural language processes, computer vision, as well as Internet of things and robotics. Hence, in this respect, AI is an umbrella term that gathers a bucket of different aspects. We can somehow look at AI to be similar to a fully-functional living being, and we can establish comparisons to figure out the degree of relationship between AI and other (sub)fields. If AI and the human body are alike, it has to possess a brain, which carries out a variety of tasks and is in charge of specific functions such the language (NLP), the sight (computer vision), and so on so forth. The body is made of bones and muscles, as much as a robot is made by circuits and metals. Machine learning can be seen as specific movements, action or thoughts we develop and that we fine-tune by doing. The Internet of things (IoT) corresponds to the human senses, which is the way in which we perceive the world around us. Finally, big data is the equivalent of the food we eat and the air we breathe, i.e., the fuel that makes us tick, as well as every input we receive from the external world that is captured by our senses. It is a quite rough comparison, but it conveys a simple way on how all the terms are related to each other.

Although many other comparisons may be done, and many of them can be correct simultaneously, the choice of what kind of features a system should have to be a proper AI is still quite controversial. In my opinion, the system should be endowed with a learning structure,an interactive communication interface, and a sensorial-like input digestion. Unfortunately, this idea is not rigorous from a scientific point of view, because it would involve a series of ethical, psychological, and philosophical considerations that should be taken into account.

### ***III. 3 Types of AI***

Instead of focusing much longer on this non-provable concept, I rather prefer to illustrate how those characteristics would reflect the different types of AI we are (and we will) dealing with.

An AI can indeed be classified in three ways: **a narrow AI**, which is nothing more than a specific domain application or task that gets better by ingesting further data and “learns” how to reduce the output error. An example here is DeepBlue for the chess game, but more generally this group includes all the functional technologies that serve a specific purpose. These systems are usually quite controllable because limited to specific tasks.

When a program is instead not programmed for completing a specific task, but it could eventually learn from an application and apply the same bucket of knowledge to different environments, we face an **Artificial General Intelligence (AGI)**. This is not *technology-as-a-service* as in the narrow case, but rather *technology-as-a-product*.

The best example for this subgroup is Google DeepMind, although it is not a real AGI in all respects. We are indeed not there yet because even DeepMind cannot perform an intellectual task as a human would. In order to get there, much more progress on the brain structure functioning, brain processes optimization, and portable computing power development have to be made.

Someone might think that an AGI can be easily achieved by piling up many narrow AIs, but in fact, this is not true: it is not a matter of number of specific skills a program can carry on, but rather the integration between all those abilities. This type of intelligence does not require an expert to work or to be tuned, as it would be the case for narrow AI, but it has a huge limitation: at the current state, it can be reached only through continuously streaming an infinite flow of data into the engine.

The final stage is instead called **Superintelligent AI (ASI)**: this intelligence exceeds largely the human one, and it is able of scientific and creative thinking; it is characterized by general common wisdom; it has social skills and maybe an emotional intelligence. Although we often assume this intelligence to be a single super computer, it is more likely that it is going to be made by a network or a swarm of several intelligences.

The way in which we will reach the different stages is though still controversial, and many schools of thoughts exist. The symbolic approach claims that all the knowledge is symbolic and the representation space is limited, so everything should be stated in formal mathematical language. This approach has historically analyzed the complexity of the real world, and it had suffered at the same time from computational problems as well as understanding the origination of the knowledge itself. The statistical AI instead focuses on managing the uncertainty in the real world (Domingos et al., 2006), which lives in the inference realm contrarily to the more deductive logical AI. On a side then, it is not clear yet to what degree the human brain should be taken as an example: biological neural network seems to provide a great infrastructure for developing an AI, especially regarding the use of sparse distributed representations (SDRs) to process information.

## *How Does AI Compare to Humans?*

So the natural question everyone is asking is “where machines stand with respect to humans?” Well, the reality is that we are still far from the point in which a superintelligence will exceed human intelligence—the so-called Singularity (Vinge, 1993). The famous futurist Raymond Kurzweil proposed in 1999 the idea of the law of accelerating returns, which envisages an exponential technological rate of change due to falling costs of chips and their increasing computational capacity. In his view, the human progress is S-shaped with inflection points corresponding to the most relevant technological advancements, and thus proceeds by jumps instead of being a smooth and uniform progress.

Kurzweil also borrowed Moore’s law to estimate accurately the precise year of the singularity: our brain is able of $10^{16}$ calculations per second (cps) and $10^{13}$ bits of memory, and assuming Moore’s law to hold, Kurzweil computed we will reach an AGI with those capabilities in 2030, and the singularity in 2045.

I believe though this is a quite optimistic view because the intelligence the machines are provided with nowadays is still only partial. They do not possess any common sense, they do not have any sense of what an object is, they do not have any earlier memory of failed attempts, they are not conscious - the so-called the “Chinese room” argument, i.e., even if a machine can perfectly translate Chinese to English and vice versa, it does not really understand the content of the conversation.

On the other side, they solve problems through structured thinking, they have more storage and reliable memory, and raw computational power. Humans instead tried to be more efficient and select ex-ante data that could be relevant (at the risk of losing some important information), they are creative and innovative, and extrapolate essential information better and faster from only a few instances, and they can transfer and apply that knowledge to unknown cases.

## References

Bostrom, N. (2014). Superintelligence: Paths, Dangers, Strategies, OUP Oxford.

Brennan, T. J., Lo, A. W. (2011). "The Origin of Behavior". Quarterly Journal of Finance, 7: 1043–1050.

Brennan, T. J., Lo, A. W. (2012). "An Evolutionary Model of Bounded Rationality and Intelligence". PLoS ONE 7(11), e50310.

Domingos, P., Kok, S., Poon, H., Richardson, M., Singla, P. (2006). "Unifying logical and statistical AI". Proceeding of the 21st National Conference on Artificial Intelligence, 1: 2–7.

Lo, A. W. (2012). "Adaptive Markets and the New World Order". Financial Analysts Journal, 68(2): 18–29.

Lo, A. W. (2013). "The Origin of Bounded Rationality and Intelligence". Proceedings of the American Philosophical Society, 157(3): 269–280.

Müller, V. C., Bostrom, N. (2016). "Future progress in artificial intelligence: A Survey of Expert Opinion", in Vincent C. Müller (ed.): Fundamental Issues of Artificial Intelligence, Springer: 553–571.

Vinge, V. (1993). "The Coming Technological Singularity: How to Survive in the Post-Human Era". In NASA. Lewis Research Center, Vision 21: Interdisciplinary Science and Engineering in the Era of Cyberspace: 11–22.

*This article is an excerpt of my book "Artificial intelligence and exponential technologies: business models evolution and new investment opportunities", edited by Springer.*

## About the Data Science Foundation

The Data Science Foundation is a professional body representing the interests of the Data Science Industry. Its membership consists of suppliers who offer a range of big data analytical and technical services and companies and individuals with an interest in the commercial advantages that can be gained from big data. The organisation aims to raise the profile of this developing industry, to educate people about the benefits of knowledge based decision making and to encourage firms to start using big data techniques.

## Contact Data Science Foundation

Email:admin@datascience.foundation
Telephone: 0161 926 3641
Atlantic Business Centre
Atlantic Street
Altrincham
WA14 5NQ
web: www.datascience.foundation